

# CSL overview for ACEs

Ben Kilminster\*, Jedong Lee\*, Kevin McFarland, Tony Vaiciulis\*  
University of Rochester

Hiroiyuki Matsunaga, Makoto Shimojima, Hirokazu Kobayashi\*  
University of Tsukuba

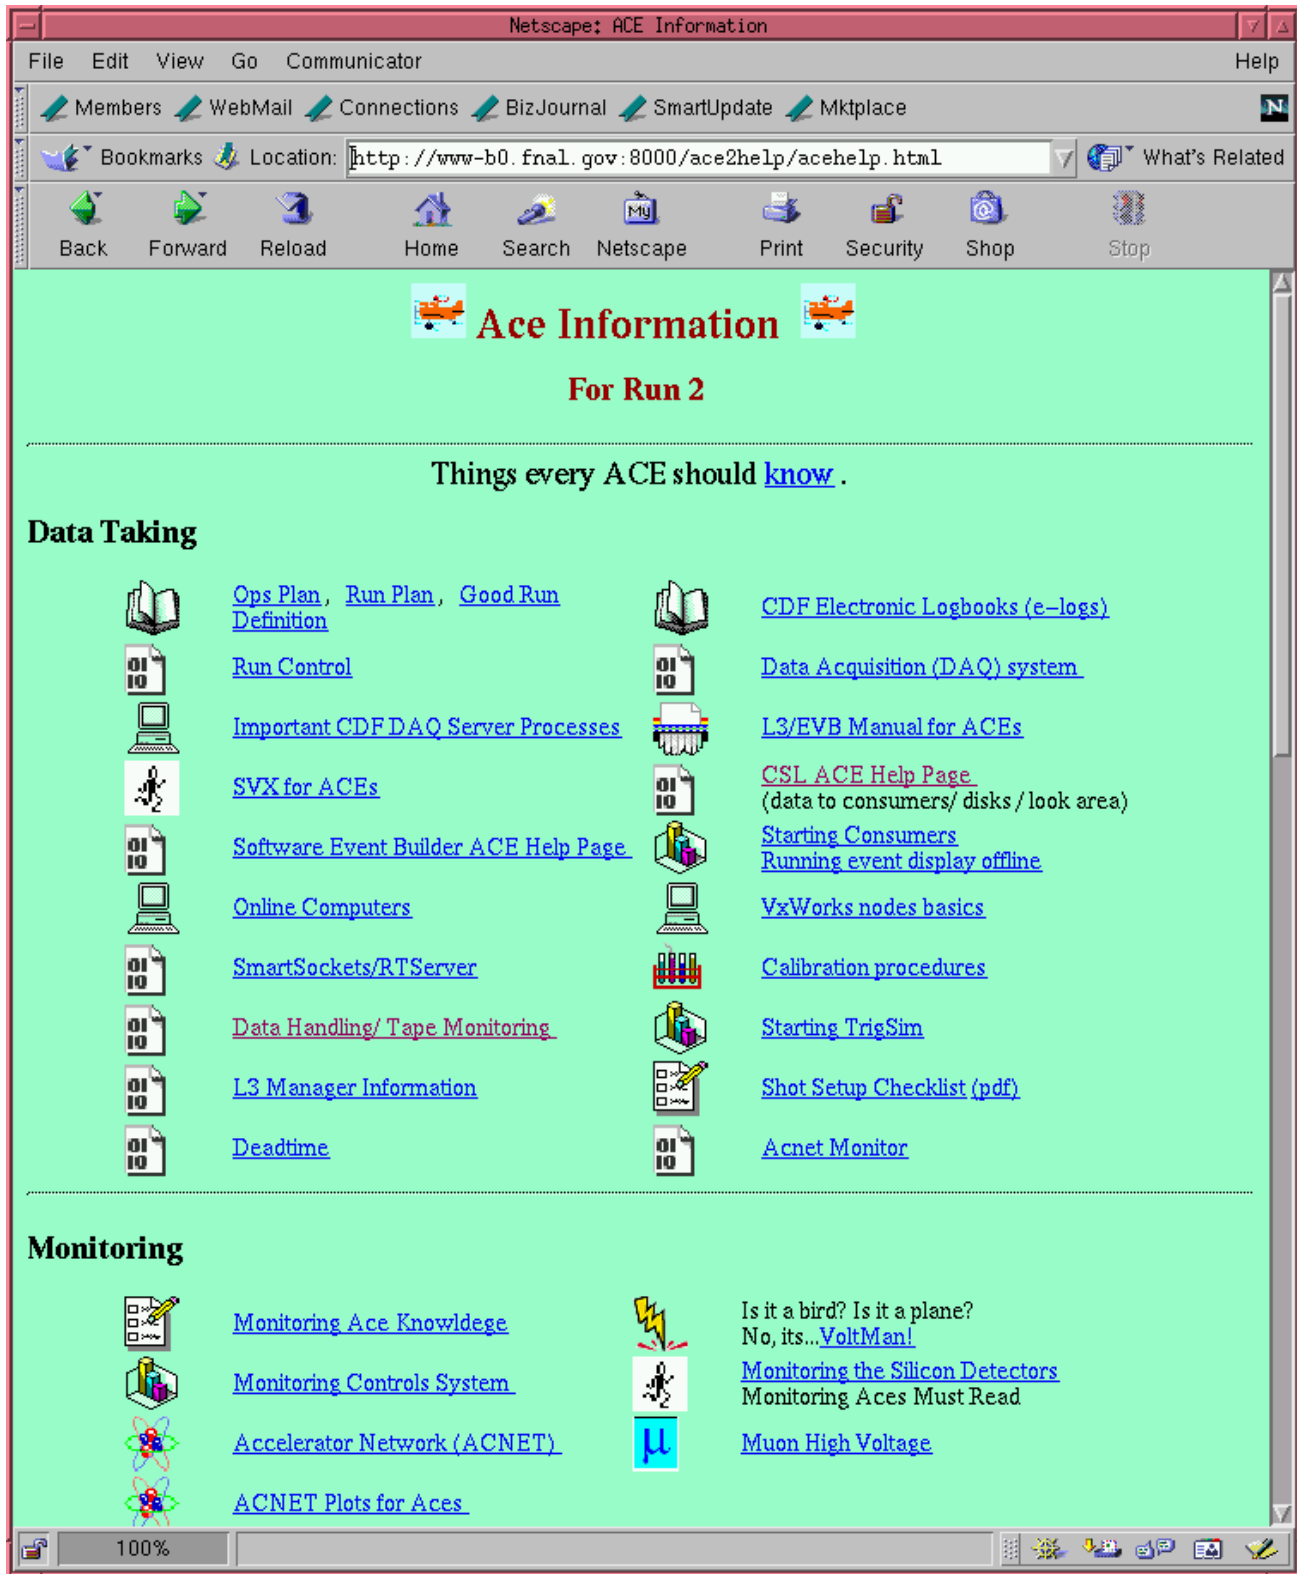
January 22, 2003

Alternates pager

- CSL description
- starting/stopping the CSL
- monitoring
- troubleshooting

# CSL ACE Help Page

Most items in this talk are documented on CSL ACE help page



# CSL ACE Help Page

Netscape: Consumer-Server/Logger ACE Help Page

File Edit View Go Communicator Help

Members WebMail Connections BizJournal SmartUpdate Mktplace

Bookmarks Location: <http://www-b0.fnal.gov:8000/ace2help/csl/> What's Related

Back Forward Reload Home Search Netscape Print Security Shop Stop

## Consumer-Server/Logger ACE Help Page

### About the CSL

Questions, comments, suggestions? Send mail to [cdf\\_csl@fnal.gov](mailto:cdf_csl@fnal.gov)

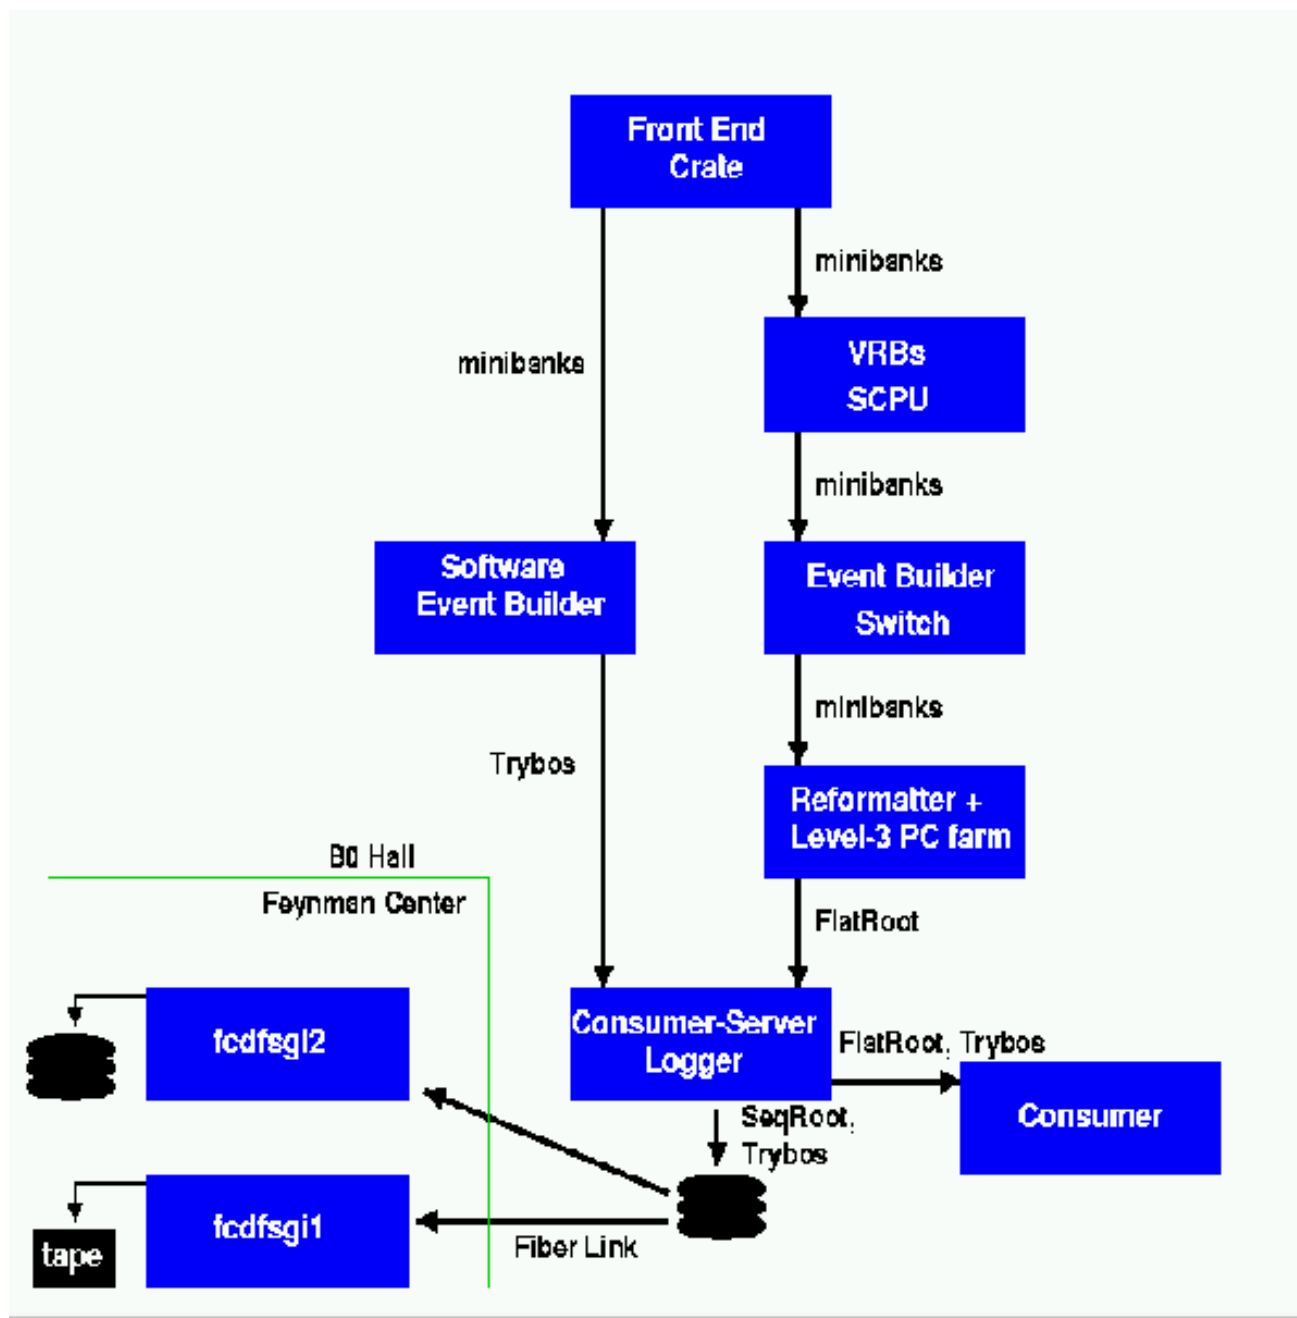
The Consumer-Server/Logger (CSL) is the central online hub through which all CDF data taken during Run II must pass. Events passing the Level-3 trigger are written to disk and served to consumers by the CSL. The raw data files on disk are then moved to the Feynman Computing Center (FCC) for further processing. Consumers are online monitoring processes, spying on a fraction of the data passing through the CSL. The diagram below shows how the CSL fits into the overall dataflow.

```
graph TD
    FEC[Front End Crate] -- minibanks --> SEB[Software Event Builder]
    FEC -- minibanks --> VRB[VRBs SCPU]
    VRB -- minibanks --> EBS[Event Builder Switch]
    EBS -- minibanks --> RLF[Reformatter + Level-3 PC farm]
    RLF -- FlatRoot --> CSL[Consumer-Server Logger]
    SEB -- FlatRoot --> CSL
    CSL -- FlatRoot --> FCD[fcdfsgl2]
    FCD -- FlatRoot --> FCD2[fcdfsgl2]
```

The diagram illustrates the dataflow architecture of the Consumer-Server/Logger (CSL). It shows the flow of data from the Front End Crate through various processing stages to the Consumer-Server Logger and finally to the Feynman Center for further processing.

**Front End Crate** connects to **Software Event Builder** and **VRBs SCPU** via **minibanks**. **VRBs SCPU** connects to **Event Builder Switch** via **minibanks**. **Event Builder Switch** connects to **Reformatter + Level-3 PC farm** via **minibanks**. **Reformatter + Level-3 PC farm** connects to **Consumer-Server Logger** via **FlatRoot**. **Software Event Builder** also connects to **Consumer-Server Logger** via **FlatRoot**. **Consumer-Server Logger** connects to **fcdfsgl2** via **FlatRoot**. **fcdfsgl2** is located in the **BO Hall Feynman Center**.

# CSL description



## primary CSL functions

- receive events from Level-3 PC farm at 20 MB/sec ( $75 \text{ Hz} \times 250 \text{ kB/event}$ )
- receive events from the software event builder
- write events to disk at 20 MB/sec
- handle as many consumer requests as possible (5-10 MB/s total)

## Starting/stopping the CSL

Page a CSL expert before starting or stopping the CSL.

During normal running, the CSL never needs to be started or stopped. It is always "on", ready to receive events.

Should it be advised by an expert :

The shift crew is able to start and stop the CSL.

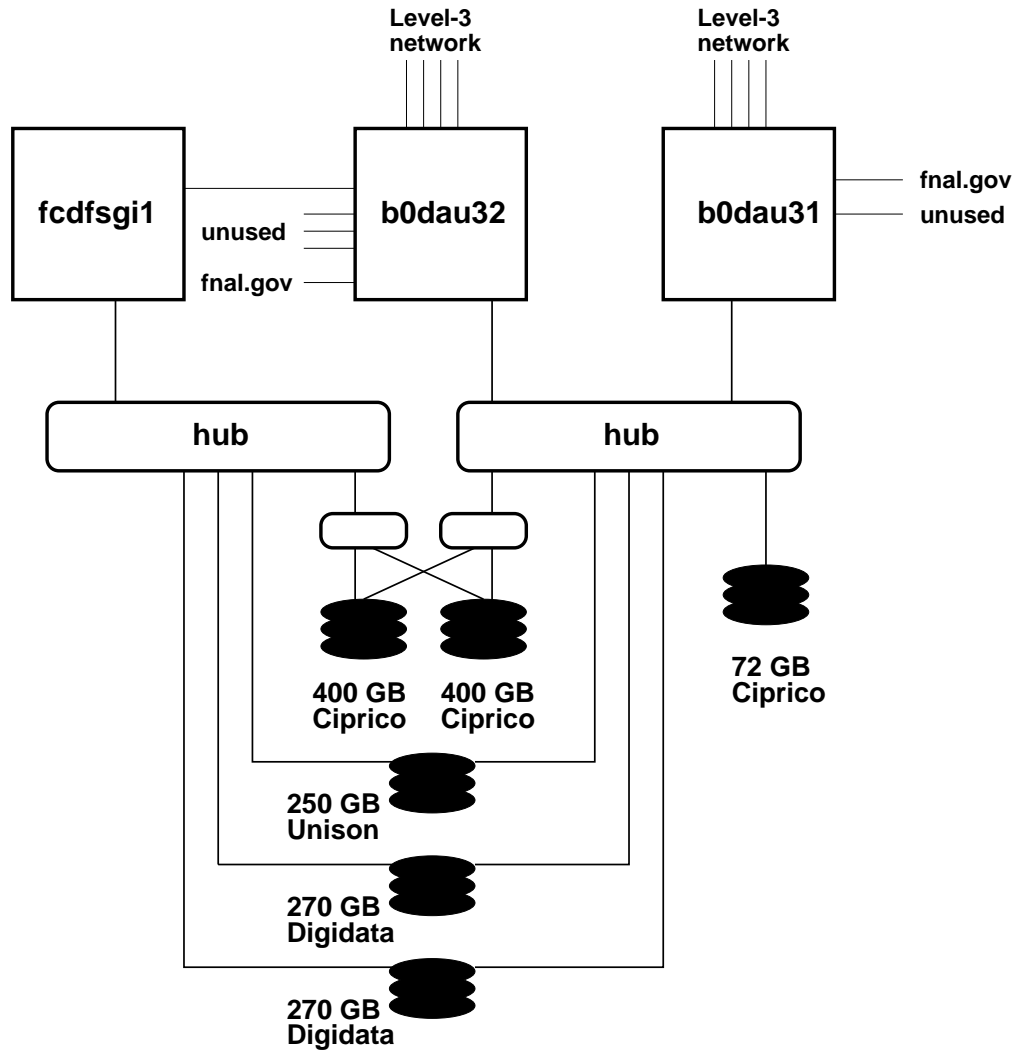
Log on to b0dau32 as user cdfdaq and type the following commands:

b0dau32: cslcom start (start all CSL processes)

b0dau32: cslcom stop (stop all CSL processes)

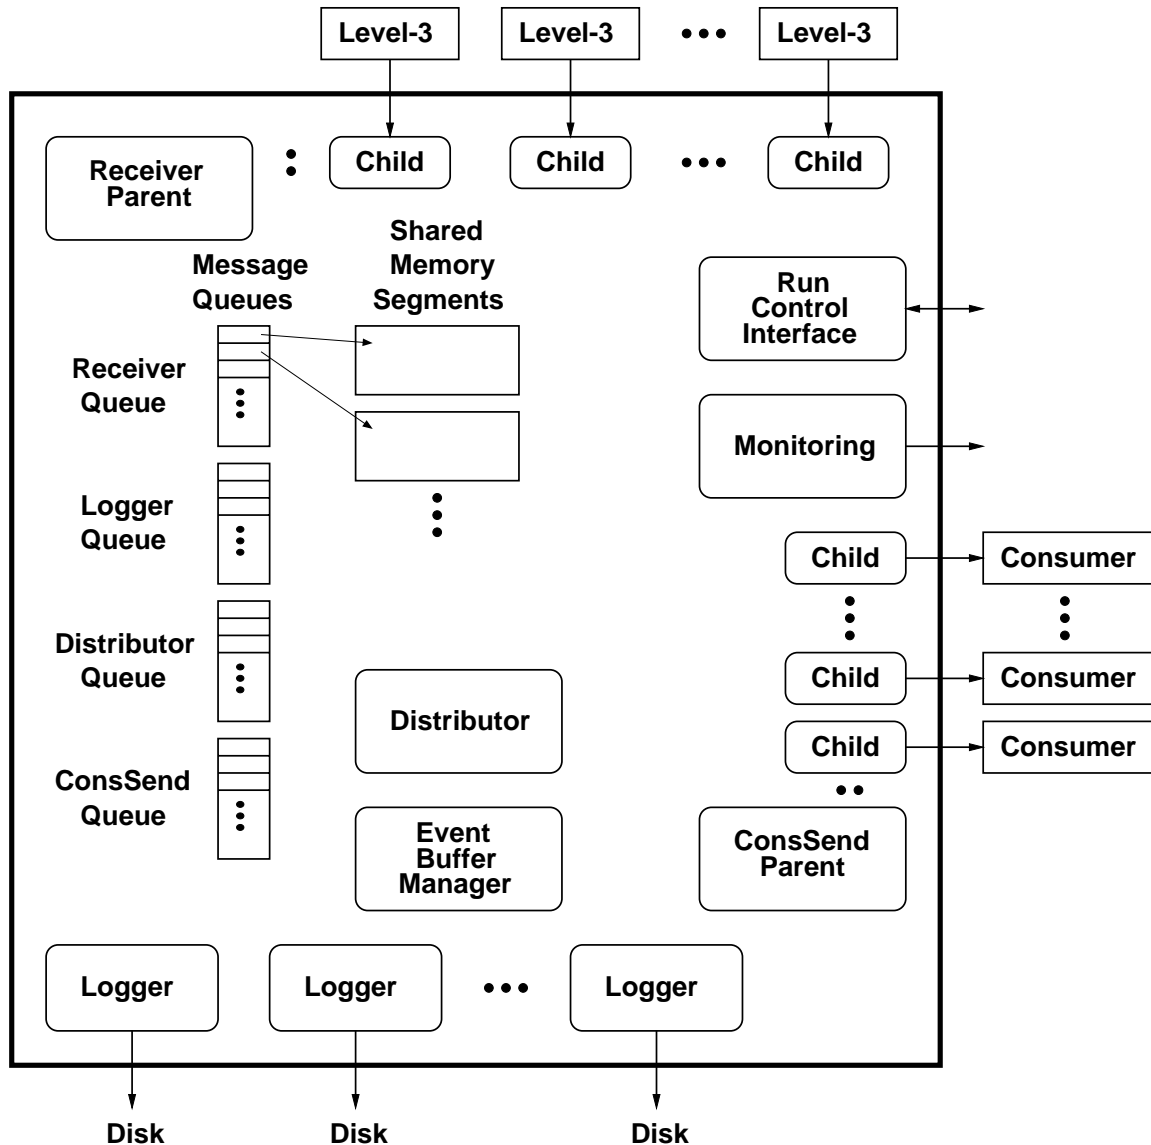
b0dau32: cslcom check (check if CSL processes are running)

# CSL hardware overview



- **b0dau32**, an SGI 2200 Server ( 4 CPUs, large I/O bandwidth) located on third floor of B0
- 1500 GB of disk space on third floor of B0 (5 RAIDs)
- dual ported disks allow both **fcdfsgi1** in FCC and **b0dau32** to access disks

# CSL software: overview



The CSL is a “server”. Possible “clients” include

- Level-3 output node processes
- software event builder processes
- 24 hour sender in partition 14
- consumers

Events are stored in shared memory buffers. Flow of events between processes inside CSL achieved by means of message queues.

# CSL software: monitoring

The monitoring process collects CSL status information and sends it via a smartsockets message to the CSL monitoring display.

There are two kinds of CSL monitoring:

- the CSL display (snapshot of current CSL state)
- the CSL history plots (show rates, disk space, events logged, etc. versus time)

Using CSL display to check current CSL status:

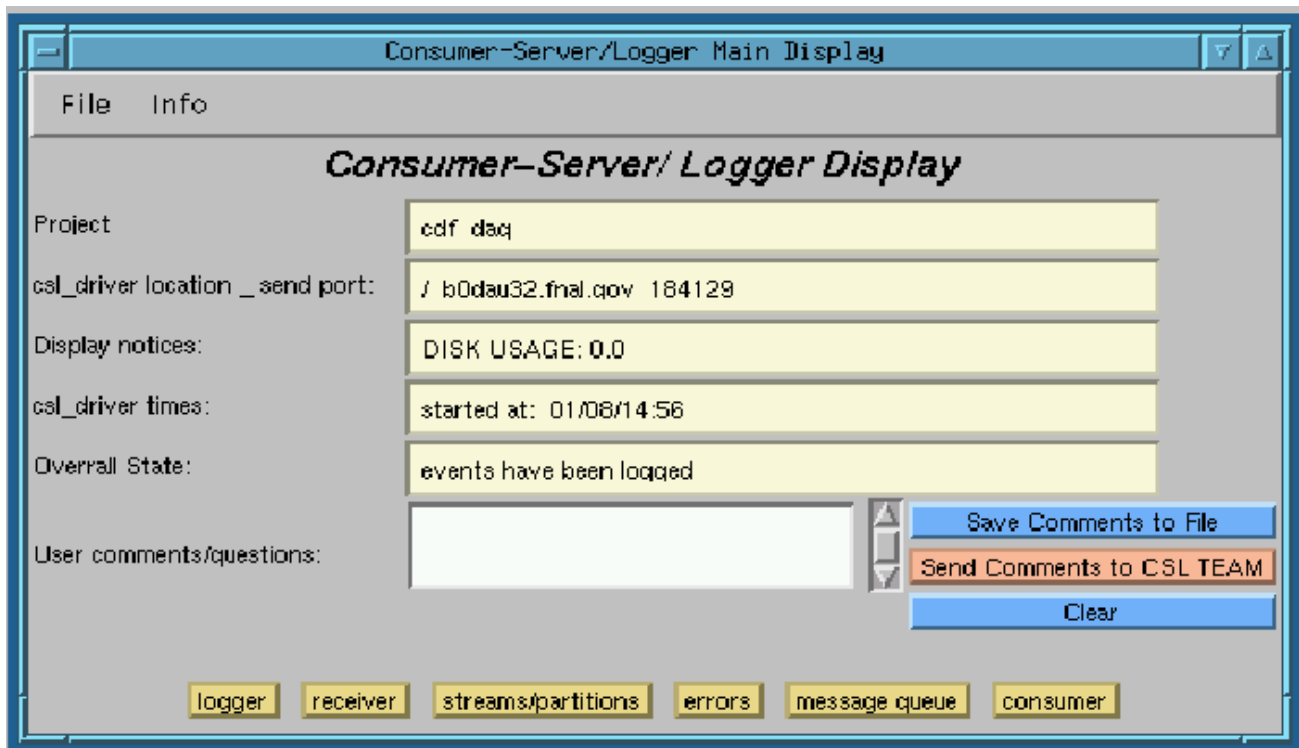
login to machine on online cluster

setup fer devel

daqmon (or cslmon for just the csl display)

A GUI should appear on your screen. Click on CSL button.

Use project name cdfdaq (this should be the default). The CSL monitor main display window should appear.





## CSL software: receiving

RECEIVER DISPLAY OF EVENTS COMING INTO CSL

FileInfo

Events From		Statistics		Status				
Client Name	Partit.	Port #	Avg Rate	Event count	Unconn	Conn	Idle	Active
b0dap24.fnal.gov	14	1762	0.0297	58				
b0dap30.fnal.gov	7	3358	0.1554	308018				
b0l3u01c	0	3018	0.146	1616				
b0l3u02c	0	2043	0.1463	1617				
b0l3u03d	0	2016	0.1468	1616				
b0l3u04d	0	2016	0.0736	808				

For each client sending events to the CSL, there is a receiver process.


The CSL monitoring GUI displays information for each receiver process: client node name, partition, number of events received, and average and instant rates (event rate, event size, activity rate).

# CSL software: logging

Use this to check whether events are being written to disk.  
For each partition and each stream this display shows


- the number of events written to disk OR
- the rate of writing events to disk in MB/sec
- CSL writes events into 1 GB output files onto disks
- CSL disks are temporary holding space for events
- After closing, file sent to FCC where fcdfsi1 puts it onto tape
- Should be enough disk space for an 8 hour buffer Some files also copied to “look” area on fcdfsi2 setup cdfsoft2 development; findfile 113453

# Look Area



- Default is to copy all Stream A physics files, the first few files from some other physics, test, cosmic, and calibration streams.
- Ace Controls
  - Aces can change from default to ALL or NONE under advice of expert
  - However ... be careful about low bandwidth for file copying and swamping look area with crap

## Paths, datasets, streams in Run 2



**path:** AND of Level-1, Level-2, Level-3 triggers.

**dataset:** OR of all paths defined for that dataset.

**stream:** collection of datasets

Events coming out of Level-3 are “streamed”: tagged as belonging to a particular stream.

- CDF DAQ system can run in multi-partition mode
- each partition is independent of other partitions
- CSL writes events into separate files for different partitions
- CSL writes events corresponding to different streams within a partition into separate files

Run II: 10 streams, 50-100 datasets

## CSL software: serving consumers

The screenshot shows a software window titled "Consumer display". It contains a table with columns: Cons. Type, Cons. id : part. id, Host name or address, Evts Req., Events Recvd, Inst Rate, Conn, Waiting, Active, and Complet.

Cons. Type	Cons. id : part. id	Host name or address	Evts Req.	Events Recvd	Inst Rate	Conn	Waiting	Active	Complet
TrigMcr	7902 / C	kCoap59 final cov	1	0	0				
(MB/s) for Current Consumers (C.C.)	C								
Current Events Requested for C.C.	1								
Combined Events Requested for C.C.	1								
Events Sent to C.C.	C								
Events Sent to ALL Consumers									

Below the table, there are two yellow buttons: "Show average rate" and "Show most recent rate".

## CSL LAW #1: write ALL events to disk

CSL LAW #2: send events to consumers as long as it does not break LAW #1







Consumers do not see all events (maybe 5-10 MB/sec). Consumers can request events by partition number, stream, L1/L2/L3 triggers.

# CSL history plots

## Consumer–Server/Logger Monitoring Plots

CSL specific plots		
receiver, logger, consumer rates	<a href="#">24 hours</a>	<a href="#">one week</a>
number of events logged	<a href="#">24 hours and one week</a>	
number of processes	<a href="#">24 hours</a>	<a href="#">one week</a>
CPU usage: receivers,loggers,consends	<a href="#">24 hours</a>	<a href="#">one week</a>
CPU usage: driver,distributor,bufman,monsend	<a href="#">24 hours</a>	<a href="#">one week</a>
general b0dau32 plots		
disk space usage	<a href="#">24 hours</a>	<a href="#">one week</a>
load average	<a href="#">24 hours and one week</a>	
free memory	<a href="#">24 hours and one week</a>	
CPU usage: global	<a href="#">24 hours</a>	<a href="#">one week</a>

# CSL history plots - disk space



## Calibration CSL

A special version of the CSL software runs on b0dap60. Useful for some calibration runs which require

- guaranteed delivery of all events to the consumer OR
- a very large event size  
(expected Run 2 event size is about 250 kB, official CSL can accept up to 3 MB, calibration CSL can accept up to 15 MB)

Calibration CSL does not log any data to disk.

## Data File Catalog

CSL writes information into the Data File Catalog database for each output data file:

- file size
- total number of events in the file
- run number
- first event number in the file
- last event number in the file
- run section numbers

This must be done before the files are put onto tape.



# Troubleshooting

Here are some reasons why you might think there is a problem with the CSL

- The CSL does not acknowledge a run control transition.
- There is an error message from the CSL in the Error Logger.
- You believe Level-3 or the software event builder is sending events to the CSL and one or more of the following is true
  - the consumers are not receiving any events
  - the events do not seem to be written to disk by the CSL
  - no files for the runs you are taking appear in the "look" area on fcdfsgi2

Some things you should check before paging a CSL expert:

- Is the CSL receiving any events?
- If the CSL is not receiving events...
- Is the CSL sending events to consumers?
- Is the CSL writing events to disk?
- Did the CSL send an error message to the Error Logger?

CSL ACE web pages describes how to answer these questions.

AFTER going through the checklist, if you still think there is a problem with the CSL, then page a CSL expert: Ben Kilminster, Hiro Matsunaga, Mako Shimojima, Tony Vaiciulis

## Recent

- Ace can :
  - Select files to go/not go to look area (LookArea menu in Run Settings)
  - Select special runs go/not go to tape (TapeOption menu in Run Settings)
- Some recent problems
  - Tape writing job fails (not CSL's fault)
    - CSL disks fill up, can be seen in CSL disk monitor display
    - Will be added to PROCMON for faster response
    - Solution : Contact Data Handling Group
  - Level 3 output nodes are all green, implying CSL is slow.
    - Sometimes happens when cosmic trigger being used during beam studies
    - Solution : Change to a beam studies trigger, or raise prescale factor
  - CSL warns that file size limit has been reached, no data will be logged
    - Happens during tests, when too much data is in a run section
    - Solution: Change trigger table, or consult with whoever wanted test, or stop run when CSL reports that file size is above 100 %.
  - CSL display exits when window closed
    - Solution: Use File → Quit instead of the “X” button